Constructing Laplace Operator from Point Clouds in $\mathbb{R}^d$ *

Mikhail Belkin † Jian Sun ‡ Yusu Wang†

Abstract

We present an algorithm for approximating the Laplace-Beltrami operator from an arbitrary point cloud obtained from a k-dimensional manifold embedded in the d-dimensional space. We show that this *PCD Laplace (Point-Cloud Data Laplace) operator* converges to the Laplace-Beltrami operator on the underlying manifold as the point cloud becomes denser. Unlike the previous work, we do not assume that the data samples are independent identically distributed from a probability distribution and do not require a global mesh. The resulting algorithm is easy to implement. We present experimental results indicating that even for point sets sampled from a uniform distribution, PCD Laplace converges faster than the weighted graph Laplacian. We also show that using our PCD Laplacian we can directly estimate certain geometric invariants, such as manifold area.

1 Introduction

The problem of analyzing and inferring underlying structure from collections of data samples is ubiquitous in various problems of science and engineering. In many cases, the input data resides or is thought to reside on or near a low-dimensional submanifold in a higher dimensional ambient space. A standard example of such data is images of an object taken under fixed lighting conditions with a rotating camera, where the dimension of the ambient space is the number of pixels, while the intrinsic dimensionality of the manifold is just two. The problem of recovering the low-dimensional manifold structure from data samples has attracted a large amount of attention recently, including the study of manifold reconstruction in the field of computational geometry, and of manifold learning in the machine learning community.

A popular class of manifold learning methods makes use of the Laplace-Beltrami operator associated to the underlying manifold. The Laplace-Beltrami operator is a fundamental geometric object and has many properties useful for practical applications. Eigenfunctions of the Laplacian form a natural basis for square integrable functions on the manifold analogous to Fourier harmonics for functions on a circle (i.e. periodic functions). Thus Laplace eigenfunctions allows one to construct a basis reflecting the geometry of the manifold, which were used to perform various tasks like dimensionality reduction [4], motion tracking [16], and intrinsic symmetry detection [19], and other applications in machine learning, computer graphics and other areas [15, 24, 25, 21, 23]. In addition, eigenvalues of the Laplace operator form the spectrum of the manifold, and can be used to estimate many important quantities of the manifold, such as manifold surface volumn and total scalar curvature. Many geometric invariants of the manifold can be reconstructed from its Laplace-Beltrami operator. For example, for a 2-manifold in $\mathbb{R}^3$, its mean curvature flow can be computed by applying the Laplace-Beltrami operator to the coordinate x, y, z, considered as functions on the manifold [11]. Finally, the Laplace operator is intimately related to heat diffusion on manifolds, and is connected to a large body of classical mathematics, connecting geometry of a manifold to the properties of its heat flow (see, e.g., [22]).

However, in most of the applications, the underlying manifold is not accessible. Instead, we are typically given a set of points from a manifold obtained through some sampling process. For example, in computer graphics, an object may be digitized by using 3-D scanning equipment. In many machine learning problems, we may reasonably assume that independent identically distributed (iid) samples are drawn from some underlying unknown probability distribution. Thus, in order to take advantage of the properties of the Laplace-Beltrami operator, one needs to build a faithful approximation of the Laplace-Beltrami operator on the underlying object from the (given) point cloud.

In this paper, we propose and implement an algorithm for provably reconstructing the Laplace-Beltrami operator in arbitrary dimensions from point cloud data. Unlike the previous work, we do not need to construct a global mesh for the data (which is challenging in high dimensions, even when the manifold is low-dimensional), or assume that the data is sampled at random iid from a probability distribution (which may not be the case in many applications such as object scanning, where sampling is obtained by a certain deterministic physical process and the samples are not independent).

*Supported by DOE under grant DE-FG02-06ER25735, by DARPA under grant HR0011-05-1-0007 and by NSF under grants CCF-0747082, DBI-0750891, and IIS-0643916.

†Department of CSE, The Ohio State University Columbus OH 43210; *mbelkin,yusu@cse.ohio-state.edu*

‡Computer Science Department, Stanford University Palo Alto CA 94305; *sunjian@stanford.edu*

Prior work. For surfaces in $\mathbb{R}^3$, several discretizations of the Laplacian have been proposed from their approximating meshes [10, 17, 20, 21, 26, 27]. One of the most popular methods is the cotangent scheme [20], which is a form of the finite element method applied to the Laplace operator on a surface. This scheme has many nice properties, including being rather simple to compute and can approximate the mean curvature flow. However, while convergence can be established for special classes of meshes [27], it is shown in [28, 29] that the cotangent scheme does not provide convergent approximation in general. In [14], the weak convergence is established for this scheme, which however does not imply convergence of the Laplacian acting on an arbitrary fixed function, nor for the Laplacian eigenfunctions. It is possible to extend the cotangent scheme to higher dimensional meshes by using the discrete Dirichlet energy, but it is not yet clear how to build it from point clouds.

Recently, Belkin *et al.* [7] proposed a new discrete scheme with convergence guarantees for arbitrary meshes. It is easy to extend their scheme to high dimensional meshes. However, although there are efficient algorithms [2, 3] for converting a point cloud from surface in $\mathbb{R}^3$ into a mesh, the mesh construction problem is rather expensive in high dimensions. The best existing such algorithms [8] take time exponential in the dimension of the ambient space, which can be much higher than the intrinsic dimension of the manifold. Our goal is to construct the Laplace operator directly from the point cloud, with running time depending only mildly on the ambient dimension.

Another line of work, originating in machine learning, studies the behavior of the weighted adjacency graph corresponding to the input points. Such graphs are of independent interest and appear in a variety of machine learning problems, including dimensionality reduction, spectral clustering and semi-supervised learning. Belkin and Niyogi [5] showed that by choosing Gaussian weights with proper bandwidth, the corresponding graph Laplacian converges to the Laplace-Beltrami operator on the manifold if the point cloud is drawn from uniform distribution. For point clouds from arbitrary probability distribution, it converges to a weighted Laplacian [5] or to manifold Laplacian up to a multiplicative factor [15].

Contributions. We present a construction of point cloud data (PCD) Laplace operator from an *arbitrary* point cloud P sampled from a k-dimensional manifold M in $\mathbb{R}^d$, where k can be much smaller than d.

1. We present the first algorithm for provable reconstruction of the Laplace operator from arbitrary point cloud data. The convergence result for our operator is established without assuming that the data is sampled iid from a probability distribution, and holds for any sufficiently dense point cloud.
2. Unlike the previous work [7], we do not need a global mesh for the data (which is difficult to obtain in high dimensions and requires complexity exponential in the dimension of the ambient space) using instead a certain local meshing procedure in the tangent space. The final complexity is linear in the ambient dimension and exponential in the intrinsic dimension.
3. We provide encouraging experimental results, showing good convergence on several test data sets. In three dimensions the results of our algorithm come to those in [7] with a global mesh. Interestingly, we show that reasonably accurate approximations of geometric invariants, specifically the surface area, can be made directly from point clouds by reconstructing the so-called *heat trace*.

The main idea behind our approach is to construct the PCD Laplace operator by building a local patch around each data point and estimating the heat kernel on each patch. Although such local patches do not form a global mesh, theoretical results of this paper show that they are sufficient to approximate the manifold Laplacian. Indeed, the experimental results show that our PCD Laplacian works well for various point clouds, and converges faster than the graph Laplacian even for points sampled from a uniform distribution. However, we point out that our algorithm requires knowing the intrinsic dimension, which is not necessary for the graph Laplacian.

2 The Algorithm

2.1 Problem Definition Consider a smooth manifold M of dimension k isometrically embedded in some Euclidean space $\mathbb{R}^d$, equipped with a natural Riemannian metric induced from the Euclidean metric. We assume that M is connected — manifolds with multiple components can be handled by applying our results in a component-wise manner. The *medial axis of* M is the closure of the set of points in $\mathbb{R}^d$ that have at least two closest points in M. For any point $w \in M$, the *local feature size at* w, denoted by $\mathrm{lfs}(w)$, is the closest distance from w to the medial axis of M. The *reach* $\rho(M)$ of M, also known as the condition number of M, is the infinum of the local feature size at any point in M. In this paper, we assume that the manifold M has a positive reach.

Given a twice continuously differentiable function $f \in C^2(M)$, let $\nabla_M f$ denote the gradient vector field of f on M. The *Laplace-Beltrami operator* Δ_M *of* f

is defined as the divergence of the gradient; that is, $\Delta_M f = \mathrm{div}(\nabla_M f)$. For example, if M is a domain in $\mathbb{R}^2$, then the Laplace operator has the familiar form $\Delta_{\mathbb{R}^2} f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$.

Given a set of points $P \in \mathbb{R}^d$, we say that P is an *ε-sampling of M* if $p \in M$ for any $p \in P$, and for any point $x \in M$, there exists $q \in P$ such that $\|x - q\| \leq \varepsilon$. The goal of this paper is to compute a *point-cloud Laplace (PCD Laplace) operator* L_P^t from an ε-sampling P of M which "approximates" Δ_M. Note that in this discrete setting, the input is a function $f : P \to \mathbb{R}$. The PCD-Laplace operator performs on f and produces another function $\mathrm{L}_P^t f : P \to \mathbb{R}$. The construction of L_P^t will be described shortly in Section 2.2. The main result of this paper is that as the sampling becomes denser, L_P^t converges to Δ_M. Note that in the theorem below, ε goes to 0 implies that $t(\varepsilon)$ also goes to 0.

Theorem 2.1 *Let the point set P_ε be an ε-sampling of the k-dimensional submanifold M of $\mathbb{R}^d$. Set $t(\varepsilon) = \varepsilon^{\frac{1}{2+\xi}}$ for an arbitrary fixed number $\xi > 0$. Then for any function $f \in C^2(M)$ we have that*

$$\lim_{\varepsilon \to 0} \sup_{P_\varepsilon} \left\| \mathrm{L}_{P_\varepsilon}^{t(\varepsilon)} f - \Delta_M f \right\|_\infty = 0,$$

where the supremum is taken over all ε-sampling of M.

In the remainder of the paper, let T_q denote the tangent space of M at point q. For a fixed point p, denote by π and Π the projection from $\mathbb{R}^d$ onto T_p and its approximation $\widehat{T}_p$ (to be introduced shortly), respectively. The *angle* between two subspaces X and Y of the same dimension in $\mathbb{R}^d$ is defined as $\angle(X, Y) = \max_{u \in X} \min_{v \in Y} \angle(u, v)$, where u and v range over all unit vectors of X and Y.

2.2 Construction of PCD Laplace Operator We now describe our reconstruction of the Laplace operator from a point cloud P sampled from M. Specifically, given a point cloud P and a parameter t, our PCD Laplace operator takes a function $f : P \to \mathbb{R}$ as **input** and produces another function $\mathrm{L}_P^t f : P \to \mathbb{R}$ as **output**. Let N be the size of P. A function f can be represented as an N-dimensional vector $[f(p_1), \ldots, f(p_N)]^T$. The output $\mathrm{L}_P^t f$ is another N-vector. Since the Laplace operator is a linear operator, the PCD Laplace operator can be represented by an $N \times N$ matrix.

Given two sets of points X and Y, let $d_H(X, Y) = \sup_{x \in X} \inf_{y \in Y} \|x - y\|$ denote the (one-sided) *Hausdorff distance* from X to Y. We assume that we know the intrinsic dimension k of the underlying manifold M. To construct the matrix L_P^t, we first describe how to compute $\mathrm{L}_P^t f(p)$ for each $p \in P$.

Algorithm PCDLaplace(P, t, p, f)

S1: Tangent space approximation. Set $r = 10t^{2+\xi}$ for any positive number ξ. Consider the set of points $P_r \subseteq P$ within distance r away from p, i.e., $P_r = P \cap B(p, r)$ where $B(p, r)$ is the d-dimensional ball centered at p with radius r. Let Q^* be the best fitting k-plane passing through p such that $d_H(P_r, Q^*)$ is minimized. Using the algorithm in Har-Peled and Varadarajan [13], we construct a 2-approximation $\widehat{T}_p$ of Q^*, i.e., $\widehat{T}_p$ is an k-plane passing through p, and $d(P_r, \widehat{T}_p) \leq 2d(P_r, Q^*)$. We return $\widehat{T}_p$ as an approximation of the tangent space T_p of M at p.

S2: Local mesh construction. Fix a constant δ. Consider the set of points P_δ that are within δ away from p, i.e., $P_\delta = P \cap B(p, \delta)$. Build the Delaunay triangulation K_δ of $\Pi(P_\delta)$ on $\widehat{T}_p$. Note that K_δ is a k-dimensional triangulation[1].

S3: Integral approximation. For a sufficiently small δ, the restriction of Π to $P_\delta \to \widehat{T}_p$ is injective. Let $\Phi : \Pi(P_\delta) \to P_\delta$ be its inverse. Let $K_{\frac{\delta}{2}}$ be the subcomplex of K_δ containing simplices whose vertices are within $\frac{\delta}{2}$ away from p. Compute $\mathrm{L}_P^t f(p)$ as:

$$\mathrm{L}_P^t f(p) = \frac{1}{(4\pi t)^{k/2} t} \sum_{\sigma \in K_{\frac{\delta}{2}}} \frac{\mathrm{A}(\sigma)}{k+1} \sum_{q \in V(\sigma)} e^{-\frac{\|p - \Phi(q)\|^2}{4t}} (f(p) - f(\Phi(q))) \quad (2.1)$$

where σ is a k-dimensional face in K_δ, $\mathrm{A}(\sigma)$ is its volume, and $V(\sigma)$ is the set of vertices of σ.

Construction of the PCD Laplace operator. Note that Eqn(2.1) is linear in f. Thus for $p_i \in P$, it can be written as $\mathrm{L}_P^t f(p_i) = R_i \mathrm{f}$, where R_i is an $N \times 1$ row vector, and $\mathrm{f} = [f(p_1), f(p_2), \ldots, f(p_N)]^T$ is the N-vector representing the input function f. Specifically, for $j \neq i$, $R_i[j] = -\frac{1}{(4\pi t)^{k/2} t} \cdot \frac{A_{p_j}}{k+1} \cdot e^{\frac{\|p_j - p_i\|^2}{4t}}$, where A_{p_j} is the volume of all k-dimensional simplices in K_δ incident to $\Pi(p_j)$; while $R_i[i]$ is the (negation of the) summation of all other elements in R_i. On the other hand, observe that R_i is simply the ith row of the matrix L_P^t. Hence performing the above procedure N times for every point in P, we obtain the PCD Laplace operator L_P^t. We remark that the resulting matrix

[1] In fact, any triangulation in $\widehat{T}_p$ works for our theoretical result, as long as the longest edge in the triangulation is bounded by $O(\varepsilon)$.

is positive semi-definite. Since each row sums up to zero, the operator is an averaging operator, and the maximum principle holds for harmonic functions of L_P^t (i.e, functions such that $\mathrm{L}_P^t f = 0$).

Finally, straightforward implementation takes $O(Nd + |P_r|^{\lceil k/2 \rceil})$ time to compute each $\mathrm{L}_P^t f(p)$, where $O(Nd)$ is the time to compute P_r. Performing it for every point in P leads to an algorithm that constructs L_P^t in $O(N^2 d + N|P_r|^{\lceil k/2 \rceil})$ time, depending exponentially on the intrinsic, instead of the extrinsic dimension.

3 Convergence of PCD Laplace Operator

In this section, we prove our main theoretical result (Theorem 2.1). Specifically, given a point cloud P that ε-samples a k-dimensional submanifold M in $\mathbb{R}^d$, we fix a point $p \in P$, and show that $\mathrm{L}_P^t f(p)$ as computed by Algorithm PCDLaplace converges to $\Delta f(p)$ for any function $f \in C^2(M)$. Let ρ be the reach of M. Recall that π and Π are the projection from $\mathbb{R}^d$ to T_p and $\widehat{T}_p$, respectively.

3.1 Approximate Tangent Space In the first step (S1) of Algorithm PCDLaplace, we compute an approximate best-fitting k-plane $\widehat{T}_p$ of the set of points $P_r = P \cap B(p, r)$, where $r = 10t^{2+\xi}$. We now show that $\widehat{T}_p$ is close to the tangent space T_p of M at p. This approach to approximate the tangent space locally was previously used in [12] for dimension detection from point clouds, and the following result was shown there.

Lemma 3.1 ([12], Lemma 6) *For any point $q \in M$ with $\|p - q\| < \rho$, we have that $\sin \angle(pq, T_p) \leq \frac{\|q-p\|}{2\rho}$, and that the distance from q to T_p is bounded by $\frac{\|q-p\|^2}{2\rho}$.*

Theorem 3.2 *Compute $\widehat{T}_p$ as in Algorithm* PCDLaplace. *$\angle(T_p, \widehat{T}_p) = O(r/\rho)$ for $r < \rho/2$ and $r \geq 10\varepsilon$.*

Proof: Recall that $\widehat{T}_p$ 2-approximates the best fitting k-plane Q^* of P_r through p, implying that $d_H(P_r, \widehat{T}_p) \leq 2d_H(P_r, Q^*) \leq 2d_H(P_r, T_p)$. This, combined with Lemma 3.1, implies that for any point $q \in P_r$, the distance from q to $\widehat{T}_p$ satisfies $d(q, \widehat{T}_p) \leq r^2/\rho$. To show that $\angle(T_p, \widehat{T}_p) = O(r/\rho)$, we show that any unit vector from T_p forms a small angle with the k-plane $\widehat{T}_p$.

First, let M_r denote the intersection between the manifold M and the ball $B(p, r - \varepsilon)$. That P is an ε-sampling of M implies that P_r is an *ε-approximation of M_r*; that is, for any point $q' \in M_r$, there is a point $q \in P_r$ such that $d(q, q') \leq \varepsilon$.

Let $B|_{T_p}(x, a)$ denote a ball in T_p, centered at x, with radius a. We claim that $B|_{T_p}(p, r/2) \subseteq \pi(M_r)$. Indeed, consider any point q on ∂M_r; $d(p, q) = r - \varepsilon$. By Lemma 3.1, $d(q, T_p) \leq (r - \varepsilon)^2/2\rho$, implying that $d(p, \pi(q)) \geq d(p, q) - d(q, T_p) \geq r/2$ when $r < \rho/2$ and $r > 10\varepsilon$. Since the projection map π is continuous, $\pi(M_r)$ contains the ball $B|_{T_p}(p, r/2)$. Note that since projection can only decrease lengths, $\pi(P_r)$ is an ε-approximation of $\pi(M_r)$, and thus an ε-approximation of (any subset of) $B|_{T_p}(p, r/2)$ as well.

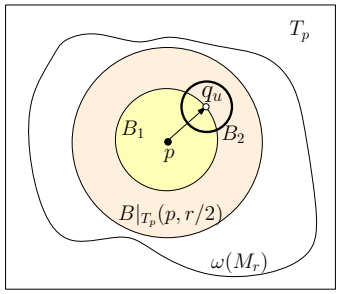

Next, consider the ball $B_1 = B|_{T_p}(p, r/4)$. For any unit direction $u \in T_p$, let q_u be the intersection point between the boundary of B_1 and the ray originated from p in direction u. See the right figure for an illustration. Let $B_2 = B|_{T_p}(q_u, 2\varepsilon)$; $B_2 \subseteq B|_{T_p}(p, r/2)$ for $r \geq 10\varepsilon$. Collect the set of points $U = \pi(P_r) \cap B_2$. We claim that q_u is contained inside the convex hull C_u of U.

Indeed, suppose that $q_u \notin C_u$. Then there exists a hyperplane H in T_p passing through q_u, such that H cuts B_2 into two halves, where one half contains C_u, and the other half does not contain any point from U and thus none from $\pi(P_r)$. This implies that there is a ball of radius ε in B_2 empty of points from $\pi(P_r)$, which impossible, as $\pi(P_r)$ is an ε-approximation of B_2. Hence $q_u \in C_u$.

Suppose that points $\tilde{q}_0, \tilde{q}_1, \ldots, \tilde{q}_k \in U$ span the k-simplex $\sigma \in C_u$ containing q_u. Let $q_i \in P_r$ be some pre-image of $\tilde{q}_i$ under the projection map π, for each $i \in [0, k]$. Since $q_i \in P_r$, we have that $\|p - q_i\| \leq r$. It then follows from Lemma 3.1 that $d(\tilde{q}_i, q_i) = d(q_i, T_p) \leq \|p - q_i\|^2/2\rho \leq r^2/2\rho$. On the other hand, we have stated at the beginning of the proof that $d(q_i, \Pi(q_i)) = d(q_i, \widehat{T}_p) \leq r^2/2\rho$. It then follows that

$$d(\tilde{q}_i, \widehat{T}_p) \;\leq\; d(\tilde{q}_i, q_i) + d(q_i, \widehat{T}_{q_i}) \;\leq\; 2r^2/\rho.$$

Since $d(q_u, \widehat{T}_p) \leq \max_{i \in [0,k]} d(\tilde{q}_i, \widehat{T}_p)$ by the convexity of the projection distance function, we have that $d(q_u, \widehat{T}_p) \leq 2r^2/\rho$. It then follows from this and that $d(p, q_u) = r/4$ (as $q_u \in \partial B_1$), that the angle between u and $\widehat{T}_p$, which is the same as the angle between pq_u and $\widehat{T}_p$, is $\arcsin \frac{d(q_u, \widehat{T}_p)}{d(p, q_u)} = O(r/\rho)$. Since this holds for any unit direction u in T_p, we conclude that $\angle(T_p, \widehat{T}_p) = \max_{u \in T_p} \angle(u, \widehat{T}_p) = O(r/\rho)$. ∎

Corollary 3.3 *For any point $q \in M$ with $\|p - q\| < \rho$, we have that $\|p - q\| = O(\|p - \Pi(q)\|)$.*

3.2 Approximate M Locally Above we established the closeness result between approximate tangent space $\widehat{T}_p$ and the tangent space T_p at a point $p \in M$. From now on, we operate on this approximate tangent space. Now consider the second step of Algorithm PCD-Laplace. Here we take the local patch of M around p, $M_\delta = M \cap B(p, \delta)$ for some fixed constant δ. The projection $\Pi : M_\delta \to \widehat{T}_p$ is injective for a sufficiently small δ and recall $\Phi : \Pi(M_\delta) \to M_\delta$ is its inverse. The main result in this section bounds the Jacobian of the map Φ. For this, we need the following two results: the first one bounds the angle between tangent spaces at nearby points on M, and the second one provides an explicit form for a certain projection map between two k-dimensional planes in $\mathbb{R}^d$.

Lemma 3.4 *For any two points $p, q \in M$ with $\|p - q\| \leq \rho/2$, we have that $cos\angle(T_p, T_q) \geq 1 - \frac{2\|p-q\|^2}{\rho^2}$.*

Proof: The proof follows from the proof of Proposition 6.2 in [18]. Specifically, modify Eqn(3) there by using the cosine-law to obtain $\cos(\theta) \geq 1 - \|w\|^2/2$ in that proof. Using the same bound for $\|w\|$ as in the proof of Proposition 6.2, together with Proposition 6.3 of [18], our claim follows. ∎

Note that our modification improves the upper bound of $cos\angle T_p, T_q$ from linear to quadratic, which is optimal and necessary for our purpose. Combined with Theorem 3.2, we obtain the following corollary saying that the Euclidean distance approximates the geodesic distance for two nearby points on a submanifold up to the third order, improving the linear order approximation given in Proposition 6.3 in [18].

Corollary 3.5 *Given two points $p, q \in M$, let $d = \|p - q\| < \rho/2$. Then we have that $d \leq d_M(p, q) \leq d + O(d^3)$.*

To bound the Jacobian of Φ, we need the following result to relate two vector spaces. The set of angles α_i's below is simply the *principal angles* between subspaces, first introduced in [1]. Different forms of this result exist. We include an elementary proof for completeness.

Lemma 3.6 *Let X and Y be two m-dimensional subspaces in $\mathbb{R}^d$ with $\angle(X, Y) < \pi/2$ (i.e., X and Y are not orthogonal). Let $h : X \to Y$ be the orthogonal projection map from $\mathbb{R}^d$ onto Y restricted on X. Then h is bijective and there exist orthonormal bases for X and Y, in which h has the following matrix form:*

$$[h] = \left(\begin{array}{ccc|c} \cos\alpha_1 & & & \\ & \ddots & & \mathbf{0} \\ & & \cos\alpha_s & \\ \hline & \mathbf{0} & & \mathbf{I}_{(m-s)\times(m-s)} \end{array}\right) \tag{3.2}$$

where $m - s = \dim(X \cap Y)$ and $0 < \alpha_s \leq \ldots \leq \alpha_2 \leq \alpha_1 = \angle(X, Y) < \pi/2$.

Proof: First h is a bijective map, as otherwise there would be a vector in X perpendicular to Y and hence $\angle X, Y = \pi/2$. Observe that without a loss of generality we can assume that $X \cap Y = \{0\}$. Indeed, if the intersection is not zero, take any orthogonal basis in the intersection. The projection matrix is identity in that basis. One proceeds by restricting X and Y to the orthogonal complement to the intersection. Set $s = m - \dim(X \cap Y)$. We now assume that X and Y are s-dimensional.

Let e_1 be the unit vector in X such that e_1 forms the maximal angle with Y. Let $f_1 = h(e_1)/\|h(e_1)\|$ and set $\alpha_1 = \angle(e_1, f_1)$; note that $h(e_1) = f_1 \cos\alpha_1$. By inductively choosing a unit vector e_k in X perpendicular to e_i for any $1 \leq i < k$ with the maximal angle with Y, and letting $f_k = h(e_k)/\|h(e_k)\|$, we obtain $\{e_i\}_{i=1}^s$ and $\{f_i\}_{i=1}^s$ such that $e_i \cdot e_j = \delta_{ij}$ and $h(e_i) = f_i \cos\alpha_i$. Obviously $\{e_i\}_{i=1}^s$ is an orthonormal basis for X.

We now show, by induction on s, that $\{f_i\}_{i=1}^s$ is an orthonormal basis for Y. Assume $f_i \cdot f_j = \delta_{ij}$ for $1 \leq i, j < k$. It suffices to show $f_k \cdot f_i = 0$ for any $1 \leq i < k$. Set $n_i = e_i - h(e_i) = e_i - f_i \cos\alpha_i$. We first show $e_k \cdot n_i = 0$ for any $1 \leq i < k$. Let $e' = (e_k \cdot n_i)e_k + (e_i \cdot n_i)e_i$. If $e' = 0$, we are done since $e_k \perp e_i$. Otherwise consider the unit vector $u = e'/\|e'\|$, which is in the orthogonal complement of $span(e_1, \cdots, e_{i-1})$ in X. By construction of e_i, u forms an angle with Y no bigger than e_i does, which means $|u \cdot n_i| \leq |e_i \cdot n_i|$, forcing $e_k \cdot n_i = 0$. Thus $f_i \cdot e_k = 0$ since f_i is the linear combination of e_i and n_i. We also have $f_i \cdot n_k = 0$ since $n_k \perp Y$. Hence $f_i \cdot f_k = 0$ as f_k is the linear combination of e_k and n_k.

By the definition of the angle between any two subspaces in $\mathbb{R}^d$, $\alpha_1 = \angle X, Y < \pi/2$. From the construction, $\alpha_1 \geq \cdots \geq \alpha_s$. The lemma follows since $h(e_i) = f_i \cos\alpha_i$ for $1 \leq i \leq s$. ∎

Theorem 3.7 *For any $x \in \Pi(M_\delta) \subset \widehat{T}_p$, there exist orthonormal bases of $\widehat{T}_p$ and $T_{\Phi(x)}$, such that*

$$D\Phi|_x = \left(\begin{array}{ccc|c} \frac{1}{\cos\alpha_1} & & & \\ & \ddots & & \mathbf{0} \\ & & \frac{1}{\cos\alpha_s} & \\ \hline & \mathbf{0} & & \mathbf{I}_{(m-s)\times(m-s)} \end{array}\right) \tag{3.3}$$

where $\angle(\widehat{T}_p, T_{\Phi(x)}) = \alpha_1 \geq \alpha_2 \geq \cdot \geq \alpha_s \geq 0$ and $m - s = \dim(\widehat{T}_p \cap T_{\phi(x)})$. The Jacobian of the map Φ

at any point $x \in \Pi(M_\delta) \subset \widehat{T}_p$ is then bounded by:

$$1 \le J(\Phi)|_x \le 1 + O\left(\frac{\|\Phi(x)-p\|^2}{\rho^2} + \frac{r}{\rho}\right).$$

Proof: For each $x \in \Pi(M_\delta)$, define $\mathrm{n}(x) = \Phi(x) - x$. Since Φ is the inverse of a projection map, n is a normal vector field on $\widehat{T}_p$. It then follows that the derivative of Φ is $D\Phi = I + D\mathrm{n}$. Note that $D\Phi|_x$ is a linear map from the tangent space of $\widehat{T}_p$ at p, which coincides with $\widehat{T}_p$, to the tangent space of M_δ at $\Phi(x)$, which is $T_{\Phi(x)}$. To compute $D\Phi|_x$, we consider $D\Phi|_x = \Pi_x^{-1} \circ \Pi_x \circ D\Phi|_x$, where $\Pi_x : T_{\Phi(x)} \to \widehat{T}_p$ is the restriction of the projection map Π on $T_{\Phi(x)}$. As long as $\angle T_{\Phi(x)}, \widehat{T}_p < \pi/2$ (as guaranteed by Lemma 3.4), Π_x is bijective and its inverse Π_x^{-1} is well defined.

We now show that $\Pi_x \circ D\Phi|_x$ is in fact the identity map, and thus $D\Phi|_x = \Pi_x^{-1}$. Indeed, let $\{e_i\}_{i=1}^m$ be an orthonormal basis for $\widehat{T}_p$. We have that $D\Phi|_x(e_i) = e_i + D\mathrm{n}|_x(e_i)$, and thus the projection of $D\Phi|_x(e_i)$ on the vector e_j is:

$$\begin{aligned}\langle D\Phi|_x(e_i), e_j\rangle &= \delta_{ij} + \langle D\mathrm{n}|_x(e_i), e_j\rangle \\ &= \delta_{ij} - \langle De_j|_x(e_i), \mathrm{n}|_x\rangle.\end{aligned}$$

The second equality follows from the fact that n is a normal vector field on $\widehat{T}_p$ and thus $\mathrm{n} \cdot e_i \equiv 0$. Furthermore, since $\widehat{T}_p$ is a flat space, $De_j(e_i) \equiv 0$ and hence $D\Phi|_x(e_j) \cdot e_i = \delta_{ij}$. This implies that $\Pi_x \circ D\Phi|_x$ is the identity map. It then follows from Lemma 3.6 (where $h = \Pi_x$ in this case) that $D\Phi_x$, has the matrix form as claimed. Finally, by Lemma 3.4 and Theorem 3.2, $\alpha_1 = \angle(\widehat{T}_p, T_{\Phi(x)}) \le O(\frac{\|x-p\|}{\rho} + \frac{r}{\rho})$. Hence the Jacobian is bounded by

$$O(\frac{1}{(\cos\alpha_1)^s}) \le 1 + O\left(\frac{\|\Phi(x)-p\|^2}{\rho^2} + \frac{r}{\rho}\right).$$

■

3.3 Proof of Theorem 2.1 We start with citing the following result which was shown in [5].

Lemma 3.8 ([5]) *There exists a sufficiently small $\delta > 0$ such that for any fixed neighborhood $\mathcal{N}_p$ of p contained in M_δ,*

$$\Delta_M f(p) = \lim_{t\to 0}\int_{\mathcal{N}_p} \frac{1}{(4\pi t)^{k/2}t} e^{-\frac{\|x-p\|^2}{4t}}(f(x)-f(p))d\nu(x). \tag{3.4}$$

Intuitively, the above lemma reduces the approximation of $\Delta_M f(p)$ to within a local neighborhood $\mathcal{N}_p$ around p. To show that our PCD Laplace at p as computed by Eqn(2.1) converges to $\Delta_M f(p)$, we first prove that the integral in Eqn(3.4) can be approximated by considering only the approximate tangent space $\widehat{T}_p$.

Lemma 3.9 *Compute $\widehat{T}_p$ as in Algorithm* PCDLaplace. *For a sufficiently small $\delta > 0$ and any fixed neighborhood $\mathcal{N}_p$ of p contained in M_δ, we have that:*

$$\Delta_M f(p) = \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} \frac{1}{(4\pi t)^{k/2}t} e^{-\frac{\|\Phi(y)-p\|^2}{4t}}(f(\Phi(y)-f(p))dy. \tag{3.5}$$

Proof: Changing variable using the map $\Phi : \widehat{T}_p \to M_\delta$ for the integral in Eqn 3.4, we have that

$$\Delta_M f(p) = \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} \frac{1}{(4\pi t)^{k/2}t} e^{-\frac{\|\Phi(y)-p\|^2}{4t}}(f(\Phi(y)-f(p))J(\Phi)|_y dy. \tag{3.6}$$

Let $RHS(I)$ denote the *right-hand side* of Eqn I. To show that Eqn 3.5 holds, it suffices to show that $|RHS(3.5) - RHS(3.6)| = 0$. Let $d_M(x,z)$ denote the shortest geodesic distance between two points $x, z \in M$. It follows from Proposition 6.3 in [18] that $d_M(x,z) \le 2\|x-z\|$ for $\|x-z\| < \rho/2$. Hence given a Lipschitz function f with Lipschitz constant $\mathrm{Lip}(f)$, we have that

$$\begin{aligned}|f(\Phi(y)) - f(p)| &\le \mathrm{Lip}(f)\cdot d_M(\Phi(y),p) \\ &\le 2\mathrm{Lip}(f)\cdot\|\Phi(y)-p\|.\end{aligned}$$

Furthermore, $\|y-p\| \le \|\Phi(y)-p\|$ and $\|\Phi(y)-p\| = O(\|y-p\|)$ by Corollary 3.3. Setting $W(t) = \frac{1}{(4\pi t)^{k/2}t}$, and combined with Theorem 3.7, we have:

$$\begin{aligned}&\big|\, RHS(3.5) - RHS(3.6)\,\big| \\ =& \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} W(t)e^{-\frac{\|\Phi(y)-p\|^2}{4t}}|f(\Phi(y))-f(p)|\cdot|J(\Phi)|_y - 1|dy \\ \le& \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} W(t)e^{-\frac{\|\Phi(y)-p\|^2}{4t}}\cdot|f(\Phi(y)-f(p)|\cdot|J(\Phi)|_y - 1|dy \\ \le& \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} W(t)e^{-\frac{\|\Phi(y)-p\|^2}{4t}}\cdot O(1)\cdot\mathrm{Lip}(f)\cdot \\ &\qquad \|\Phi(y)-p)\|O(\frac{\|\Phi(y)-p\|^2}{\rho^2}+\frac{r}{\rho})dy \\ \le& \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} W(t)e^{-\frac{\|y-p\|^2}{4t}}\mathrm{Lip}(f)\|y-p\|O(\frac{\|y-p\|^2}{\rho^2}+\frac{r}{\rho})dy \\ \le& \lim_{t\to 0}\int_{\Pi(\mathcal{N}_p)} \frac{1}{(4\pi t)^{k/2}t} e^{-\frac{\|y-p\|^2}{4t}}O(\|y-p\|^3 + r\|y-p\|)dy \\ =& \lim_{t\to 0} O(\sqrt{t}) \;=\; 0.\end{aligned}$$

The last line follows from the Claim 3.1 below and from the fact that $r = O(t^{2+\xi})$ for a positive ξ. ■

Claim 3.1 *For any constant $c \geq 0$, we have that $\int_{\mathbb{R}^k} e^{-\frac{\|y\|^2}{4t}} \|y\|^c dy = O(t^{\frac{k}{2}} t^{\frac{c}{2}})$.*

Proof: Note that y is a vector in $\mathbb{R}^k$. Writing y as $y = \langle y_1, \ldots, y_k \rangle$, the input integral is the same as

$$\int_{\mathbb{R}^k} e^{-\frac{y_1^2+y_2^2+\cdots+y_k^2}{4t}} (y_1^2 + \cdots + y_k^2)^{c/2} dy_1 dy_2 \cdots dy_k.$$

Now use the substitution $z = y/\sqrt{4t}$. That is $z = \langle z_1, \ldots, z_k \rangle$, and $z_i = y_i/\sqrt{4t}$ for each $i \in [1, k]$. We have that the integral in the claim equals to:

$$\begin{aligned} & \int_{\mathbb{R}^k} e^{-\|z\|^2} \|z\|^c (4t)^{c/2} d(z\sqrt{4t}) \\ = \; & (4t)^{c/2+k/2} \int_{\mathbb{R}^k} e^{-\|z\|^2} \|z\|^c dz \;=\; O(t^{c/2+k/2}). \end{aligned}$$

■

Proof of Theorem 2.1. We now show that the integral computed in Eqn(2.1) of Step 3 approximates the right-hand side of Eqn(3.5), thereby implying that in the limit, $\mathrm{L}_P^t f(p)$ converges to $\Delta_M f(p)$ for any p. There are several parameters involved in the algorithm so far. To summarize, given an ε-sampling P of M, we choose $t = \varepsilon^{\frac{1}{2+\xi}}$. The size of the neighborhood used to approximate the tangent space $\widehat{T}_p$ is $r = 10\varepsilon = 10t^{2+\xi}$. The size of the neighborhood to construct the triangulation is a fixed constant $\delta < \rho/4$ independent of t. Easy to verify that when ε goes to 0, all the conditions in previous lemma and theorems can be satisfied[2].

Recall that K_δ is the Delaunay triangulation of the set of projected points $\Pi(P_\delta) = \Pi(P \cap B(p, \delta))$ in $\widehat{T}_p$, and $K_{\delta/2}$ contains triangles whose vertices are within $\delta/2$ distance to the point p. We claim that each simplex $\sigma \in K_{\frac{\delta}{2}}$ has diameter at most 2ε. This is because that since P is an ε-sampling of M, it can be verified that the Voronoi cell of any vertex $q \in K_{\frac{\delta}{2}}$ is contained in the k-ball in $\widehat{T}_p$ centered at q with radius ε. As such, the length of any edge in σ is at most 2ε. Now consider a simplex $\sigma \in K_{\frac{\delta}{2}}$, and let q be a vertex of σ. We now show that for any point $y \in \sigma$, $\|\Phi(q) - \Phi(y)\| = O(\varepsilon)$.

Indeed, let $l : [0, 1] \to qy$ parametrize the segment qy, and $\gamma = \Phi(qy)$ the image (curve) of qy on M. The length of qy is bounded by 2ε. The length of γ is $\|\gamma\| = \int_0^1 \|d(\Phi(l(t)))\| dt$. On the other hand, it follows from Theorem 3.7 that the largest eigenvalue λ_{max} of $D\Phi|_x$ is $\frac{1}{\cos\alpha_1}$, where

$$\alpha_1 = \angle(\widehat{T}_p, T_{\Phi(x)}) = O(\|\Phi(x) - p\|/\rho + r/\rho)$$

by Lemma 3.1 and Theorem 3.2. Since all points x we consider are inside $\Pi(M_\delta)$, we have that $\|\Phi(x) - p\| \leq \delta$. This implies that $\alpha_1 = O(\delta/\rho + r/\rho)$ and thus

$$\lambda_{max} = 1/\cos\alpha_1 = 1 + O(\delta/\rho + r/\rho) = O(1)$$

by using Taylor expansions. Finally, observe that

$$\|d(\Phi(l(t)))\| = \|D\Phi|_{l(t)} d(l(t))\| \leq \lambda_{max} \|d(l(t))\|.$$

It then follows that:

$$\begin{aligned} \|\Phi(q) - \Phi(y)\| \leq \|\gamma\| \;&=\; \int_0^1 \|D(\Phi(l(t)))\| \\ &\leq \int_0^1 \lambda_{max} \|dl(t)\| \;=\; \lambda_{max}\|q - y\| \;=\; O(\varepsilon). \end{aligned}$$

Combined this with the triangle inequality, we have that $\|\Phi(q) - p\|^2 - \|\Phi(y) - p\|^2 = O(\varepsilon)$. Furthermore, since f is a Lipschitz function, we have that $f(\Phi(q) = f(\Phi(y)) + O(\varepsilon)$. Putting everything together and setting $K_{\frac{\delta}{2}}$ as the neighborhood $\Pi(\mathcal{N}_p)$ in Eqn (3.5), we obtain:

$$\begin{aligned} & \lim_{t\to 0} \Big| \; \mathrm{L}_P^t f(p) - RHS(3.5) \; \Big| \\ = \; & \lim_{t\to 0} \Big| \sum_{\sigma \in K_{\frac{\delta}{2}}} \int_\sigma \frac{1}{(4\pi t)^{k/2} t} \big[\, e^{-\frac{\|\Phi(q)-p\|^2}{4t}} (f(\Phi(q)) \\ & -f(p)) - e^{-\frac{\|\Phi(y))-p\|^2}{4t}} (f(\Phi(y)) - f(p)) \, \big] dy \Big| \\ \leq \; & \lim_{t\to 0} \sum_{\sigma \in K_{\frac{\delta}{2}}} \int_\sigma \frac{1}{(4\pi t)^{k/2} t} e^{-\frac{\|y\|^2}{4t}} \cdot \big[|f(\Phi(y)) - f(p)| \\ & \cdot (e^{\frac{O(\varepsilon)}{4t}} - 1) + e^{\frac{O(\varepsilon)}{4t}} O(\varepsilon) \big] dy \\ \leq \; & \lim_{t\to 0} \frac{2\|f\|_\infty (e^{\frac{O(\varepsilon)}{4t}} - 1) + e^{\frac{O(\varepsilon)}{4t}} O(\varepsilon)}{t} \int_{K_{\frac{\delta}{2}}} \frac{e^{-\frac{\|y\|^2}{4t}} dy}{(4\pi t)^{k/2}} \\ \leq \; & \lim_{t\to 0} \frac{O(t^{1+\xi}) + O(t^{2+\xi})}{t} \;=\; \lim_{t\to 0} O(t^\xi) \;=\; 0. \end{aligned}$$

Theorem 2.1 then follows from the above inequality and Lemma 3.9.

Rate of convergence. We remark here that one can obtain the rate of convergence for our algorithm with respect to either t or to ε. In particular, it is shown in [5] that the difference between $\Delta_M f(p)$ and the integral in RHS(3.4) is bounded by $O(t^{\frac{1}{2}})$. The difference between the integral in RHS(3.4) and that in RHS(3.5) is bounded by $O(t^{1/2})$, while that between the integral in RHS(3.5) and $\mathrm{L}_P^t f(p)$ is bounded by $O(\frac{\varepsilon}{t^2})$. Thus the total error is of the order of $O(\frac{\varepsilon}{t^2}) + O(t^{1/2})$. It is easy to see that the best rate is obtained by choosing $t = \varepsilon^{\frac{2}{5}}$, which gives us the rate of $O(\varepsilon^{\frac{1}{5}})$.

[2] In practice we may not know ε. One can choose t conservatively. All results hold as long as $t \geq \varepsilon^{\frac{1}{2+\xi}}$ and t goes to 0.

4 Experiments and Applications

In this section, we compute the PCD Laplacian for different data, showing its convergence as the point cloud becomes denser, and comparing its performance with the weighted graph Laplacian and mesh Laplacian. We also present some preliminary results on manifold area estimation using the spectrum of PCD Laplacian. We note that in our implementation, we estimate tangent spaces using the Principal Component Analysis (PCA). Although it does not provide theoretical guarantees, it is easy to implement and is more robust to outliers in practice.

Experimental setup. To analyze the convergence behavior, we need the "ground truth", that is, we need to know the Laplace operator for the underlying manifold from which the point cloud is sampled. This somewhat limits the type of surfaces that we can experiment with. In this paper, we consider 2-sphere S^2, flat 2-torus T^2 and flat 3-torus T^3. The description of these manifolds is in Table 1. We obtain the point clouds that sample S^2 both uniformly and non-uniformly. To sample S^2 uniformly, we draw three coordinates (x, y, z) independently from the normal distribution and then normalizing them onto the sphere. To achieve a non-uniform sampling of S^2, we impose a distribution of Gaussian mixture and accept a sample point with higher probability if the Gaussian mixture has the bigger value at the sample point. We also obtain the point clouds that sample T^2 and T^3 uniformly and non-uniformly. To sample T^2or T^3 uniformly, we only need to uniformly sample the parameter domain (square or cube) since the parameterization is isometric. The non-uniform sampling of T^2 and T^3 are computed using the Gaussian mixture strategy as before.

Given an input function $f : M \to \mathbb{R}$ defined on a manifold M, and a point cloud P sampling M, we evaluate the manifold Laplacian and the PCD Laplacian at each point in P, and obtain two vectors U and $\widehat{U}$, respectively. To measure the error of the PCD Laplacian, we consider the commonly used normalized L_2 error $E_2 = \frac{\|U-\widehat{U}\|_2}{\|U\|_2}$. Notice that our theoretical result is that our PCD Laplacian converges under the L_∞ *norm* (i.e, point-wise convergence), which is a stronger result than the L_2*-convergence*. So we also show the normalized L_∞ error $E_\infty = \frac{\|U-\widehat{U}\|_\infty}{\|U\|_\infty}$.

Results on fixed functions. We experiment three different functions: $f = x$, $f = x^2$ and $f = e^x$. Their manifold Laplacian can be computed explicitly (Table 1). We compare the PCD Laplacian and the weighted graph Laplacian [15] with the manifold Laplacian. Table 2 shows the comparisons for uniform point clouds of the 2-sphere S^2. As we can see, both discrete Laplace operators show convergence behavior in this case. However, the PCD Laplacian converges much faster than the weighted graph Laplacian even for this uniformly sampled data. The approximation error of PCD Laplacian is more than one magnitude smaller. Note that the approximation error of PCD Laplacian is in fact comparable with that of mesh Laplacian (the error of the mesh Laplacian for spheres is reported in [7]). Similar results are observed for non-uniformly sampled points as shown in Table 3. On flat torus T^2 and T^3, we observe the similar convergence results. Due to the space limit, we only show the normalized L_2 error for the non-uniform sampling data; see Table 4.

Results on manifold area estimation. Our estimation of the manifold area is based on the so called heat trace [9]. The heat trace of a manifold $Z(t) = \sum_i e^{-\lambda_i t}$ is a spectral invariant where λ_i is the i-th eigenvalue of its Laplace-Beltrami operator. It has an asymptotic expansion as $t \to 0^+ : Z(t) = (4\pi t)^{-k/2} \sum_{i=1}^{\infty} a_i t^i$ where a_i's are integrals over M of polynomial in curvature and its covariant derivatives. In particular, we have $a_0 = Area(M)$ and $a_1 = \frac{1}{6}\int_M S$ where S is the scalar curvature. We can take $(4\pi t)^{k/2} z(t)$ as zero order approximation of $a_0 = Area(M)$. Table 5 shows the manifold area estimation via the heat trace estimated by the first 200 smallest eigenvalues.

5 Discussion

In this paper, we provide the first provable reconstruction of the Laplace-Beltrami operator from an arbitrary set of points. The time complexity of our algorithm depends exponentially on the intrinsic, rather than the extrinsic dimension of the underlying manifold. We remark that it appears possible to choose δ in our algorithm to be $t^{1-\beta}$ for any positive β. This would imply that the time complexity of our algorithm can be further improved to $O(N(\frac{1}{\varepsilon})^{O(k)})$ (that is, the dependency on the number of points is linear). This is one of the future work.

We note that our results hold if when constructing $\mathrm{L}_P^t f(p)$, we use the distance between projected points in $\widehat{T}_p$ instead of distances in $\mathbb{R}^d$ (that is, replacing $\|\Phi(q) - p\|$ in Eqn ( ) with $\|q - p\|$). This may be useful in analyzing the behavior of our algorithm for noisy point clouds, as the projected distance is somewhat more robust.

Finally, based on recent results by Belkin and Niyogi [6], eigenfunctions of our PCD Laplacian also converges to the eigenfunctions of the Laplace operator

on the underlying manifold. It thus makes it promising to use our PCD Laplacian to approximate spectral invariant quantities such as manifold area (as demonstrated in our paper) and total scalar curvature. It will be interesting to future investigate this direction.

Acknowledgment: The authors would like to thank anonymous reviewers for helpful comments, including the pointer to principal angles.

References

[1] S. N. Afriat. Orthogonal and oblique projectors and the characteristics of pairs of vector spaces. *Proc. Cambridge Philos. Soc.*, 53:800–816, 1957.

[2] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. *Discr. Comput. Geom.*, 22:481–504, 1999.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. & Applications*, 12:125–141, 2002.

[4] M. Belkin and P. Niyogi. Laplacian Eigenmaps for dimensionality reduction and data representation. *Neural Computation*, 15(6):1373–1396, 2003.

[5] M. Belkin and P. Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. In *COLT*, pages 486–500, 2005.

[6] M. Belkin and P. Niyogi. Convergence of Laplacian Eigenmaps. Preprint, 2008.

[7] M. Belkin, J. Sun, and Y. Wang. Discrete Laplace operator on meshed surfaces. In *SCG '08: Proceedings of the twenty-fourth annual symposium on Computational geometry*, pages 278–287, New York, NY, USA, 2008. ACM.

[8] J. D. Boissonnat, L. J. Guibas, and S. Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. In *Proc. 23rd ACM Sympos. on Comput. Geom.*, 2007.

[9] M. Craioveanu, M. Puta, and T. Rassias. *Old and new aspects in spectral geometry in Mathematics and applications*. Springer, 2001.

[10] L. Demanet. Painless, highly accurate discretizations of the Laplacian on a smooth manifold. Technical report, Stanford University, 2006.

[11] M. Desbrun, M. Meyer, P. Schröder, and A. H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. *Computer Graphics*, 33(Annual Conference Series):317–324, 1999.

[12] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds with high codimension. In *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, pages 329–337, New York, NY, USA, 2003. ACM.

[13] S. Har-Peled and K. R. Varadarajan. Projective clustering in high dimensions using core-sets. In *Proc. 18th Annu. ACM Sympos. Comput. Geom.*, pages 312–318, 2002.

[14] K. Hildebrandt, K. Polthier, and M. Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata*, 123(1):89–112, December 2006.

[15] S. Lafon. *Diffusion Maps and Geodesic Harmonics*. PhD. Thesis, Yale University, 2004.

[16] Z. Lu, M. C. Perpinan, and C. Sminchisescu. People Tracking with the Laplacian Eigenmaps Latent Variable Model. In *Advances in Neural Information Processing Systems*, 2007.

[17] U. F. Mayer. Numerical solutions for the surface diffusion flow in three space dimensions. *comput. Appl. Math*, 20(3):361–379, 2001.

[18] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1):419–441, 2008.

[19] M. Ovsjanikov, J. Sun, and L. J. Guibas. Global intrinsic symmetries of shapes. In *SGP to appear*, 2008.

[20] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1):15–36, 1993.

[21] M. Reuter, F.-E. Wolter, and N. Peinecke. Laplace-Beltrami spectra as "Shape-DNA" of surfaces and solids. *Computer-Aided Design*, 38(4):342–366, 2006.

[22] S. Rosenberg. *The Laplacian on a Riemannian Manifold: An Introduction to Analysis on Manifolds*. Cambridge University Press, 1997.

[23] R. M. Rustamov. Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 225–233, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.

[24] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.

[25] O. Sorkine, Y. Lipman, D. Cohen-Or, M. Alexa, C. Rössl, and H.-P. Seidel. Laplacian surface editing. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pages 179–188. ACM Press, 2004.

[26] G. Taubin. A signal processing approach to fair surface design. In *SIGGRAPH '95*, pages 351–358, New York, NY, USA, 1995. ACM Press.

[27] G. Xu. Discrete Laplace-Beltrami operators and their convergence. *Comput. Aided Geom. Des.*, 21(8):767–784, 2004.

[28] G. Xu. Convergence analysis of a discretization scheme for Gaussian curvature over triangular surfaces. *Comput. Aided Geom. Des.*, 23(2):193–207, 2006.

[29] Z. Xu, G. Xu, and J.-G. Sun. Convergence analysis of discrete differential geometry operators over surfaces. In *IMA Conference on the Mathematics of Surfaces*, pages 448–457, 2005.

Manifold	Parameterization	Laplace-Beltrami operator
S^2	$\phi : [0, 2\pi]^2 \to \mathbb{R}^3$ $(\alpha, \beta) \mapsto (\sin\alpha sin\beta, \cos\alpha \sin\beta, cos\beta)$	$\Delta_{S^2} f = \frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\big(\sin\theta\frac{\partial f}{\partial\theta}\big) + \frac{1}{\sin^2\theta}\frac{\partial^2 f}{\partial\phi^2}$
T^2	$\phi : [0, 2\pi]^2 \to \mathbb{R}^4$ $(\alpha, \beta) \mapsto (\cos\alpha, \sin\alpha, \cos\beta, \cos\beta)$	$\Delta_{T^2} f = \frac{\partial^2 f}{\partial\alpha^2} + \frac{\partial^2 f}{\partial\beta^2}$
T^3	$\phi : [0, 2\pi]^3 \to \mathbb{R}^6$ $(\alpha, \beta, \theta) \mapsto (\cos\alpha, \sin\alpha, \cos\beta, \sin\beta, \cos\theta, \sin\theta)$	$\Delta_{T^3} f = \frac{\partial^2 f}{\partial\alpha^2} + \frac{\partial^2 f}{\partial\beta^2} + \frac{\partial^2 f}{\partial\theta^2}$

Table 1: Description of S^2, T^2 and T^3.

function	1000	2000	4000	8000	16000
$f = x$	0.425 / 1.116	0.158 / 0.419	0.070 / 0.171	0.033 / 0.086	0.022 / 0.034
$f = x^2$	0.276 / 0.566	0.092 / 0.210	0.039 / 0.089	0.017 / 0.030	0.008 / 0.015
$f = e^x$	0.415 / 0.578	0.138 / 0.220	0.062 / 0.089	0.028 / 0.034	0.020 / 0.020
PCD Laplacian					
$f = x$	0.794 / 1.191	0.641 / 1.147	0.527 / 1.014	0.406 / 1.059	0.285 / 0.780
$f = x^2$	0.590 / 0.917	0.426 / 0.825	0.350 / 0.761	0.251 / 0.483	0.179 / 0.308
$f = e^x$	0.746 / 0.748	0.603 / 0.661	0.483 / 0.632	0.333 / 0.389	0.237 / 0.328
weighted graph Laplacian					

Table 2: Normalized L_2 / L_∞ error for uniform sampling on S^2.

function	1000	2000	4000	8000	16000
$f = x$	0.466 / 1.780	0.174 / 0.722	0.079 / 0.260	0.041 / 0.125	0.037 / 0.077
$f = x^2$	0.268 / 0.517	0.104 / 0.274	0.046 / 0.117	0.020 / 0.072	0.013 / 0.019
$f = e^x$	0.398 / 0.544	0.136 / 0.224	0.058 / 0.095	0.036 / 0.087	0.031 / 0.028
PCD Laplacian					
$f = x$	0.931 / 2.697	0.720 / 1.753	0.478 / 0.968	0.362 / 0.950	0.269 / 0.631
$f = x^2$	0.550 / 0.922	0.429 / 0.869	0.305 / 0.521	0.205 / 0.376	0.146 / 0.241
$f = e^x$	0.820 / 0.989	0.613 / 0.826	0.389 / 0.566	0.308 / 0.381	0.210 / 0.243
weighted graph Laplacian					

Table 3: Normalized L_2/L_∞ error for non-uniform sampling on S^2.

function	2500	7500	10000	15000	2500	7500	15000	20000
	T^2				T^3			
$f = x$	0.927	0.225	0.156	0.090	0.736	0.145	0.100	0.065
$f = x^2$	0.437	0.115	0.078	0.046	0.468	0.110	0.083	0.061
$f = e^x$	0.800	0.208	0.142	0.084	0.670	0.139	0.101	0.074
PCD Laplacian								
$f = x$	1.546	0.858	0.743	0.658	1.015	0.602	0.420	0.370
$f = x^2$	0.739	0.415	0.357	0.314	0.525	0.312	0.223	0.198
$f = e^x$	1.339	0.728	0.646	0.571	0.847	0.467	0.335	0.296
weighted graph Laplacian								

Table 4: Normalized L_2 error for non-uniform sampling on T^2 and T^3.

Model(—P—)	S^2(8000)	T^2(10000)	T^3(10000)	Eight(3070)	Genus3(26620)
EST. / TRU.	13.10 / 12.57	40.41 / 39.49	277.7 / 248.1	0.990 / 0.998	1.977 / 2.023

Table 5: Manifold area ESTimation. For the mesh model Eight and Genus3, their TRUe manifold areas are estimated by summing the areas of all the triangles .